Science Working Group for the AM Platform (SWAMP)

Report of a Workshop on EOS Data Issues held on June 01-02, 2000

August 24, 2000

Authors

(Representatives of SWAMP data processing, SWAMP science, and EOS Project)

B. Barkstrom (CERES)

F. Bordi (EOS Project)

G. Bothwell (MISR)

C. Justice (MODIS)

E. Masuoka (MODIS)

R. Wolfe (MODIS)

D. Ziskin (MOPITT)

Abstract

In response to an action from the March 2000 meeting of the Terra science working group (known as SWAMP), representatives of the Terra instrument teams, DAACs, EOS and ESDIS Projects, and the Aqua Project met in a workshop on June 1-2 with the following intent:

- to evaluate how well the EOS Data and Information System (EOSDIS) supports current and impending EOS missions
- to determine how requirements have changed with respect to system capacity
- to recommend solutions that appear needed to meet evolving needs.

The participants discussed the current operating status of EOSDIS, in particular the lower than expected throughput and how it should be addressed. They noted that the February 1996 baseline sizing used to implement EOSDIS is not adequate to support the science data needs. Because that baseline was established before the algorithms were developed and could be run in the production environment, it did not have a clear empirical basis. In addition, the 1996 baseline does not appear to have been based on previous NASA experience in validating and producing Earth science data. Terra instrument team representatives presented revised system sizing estimates based on current experience and improved understanding of the EOS production environment. The group noted that the current performance of the system has yet to meet an operational level of production equivalent to the Option A+ first year capacity of 1x of Level 1 products and .5X of Level 2 and higher products as volumes specified in the 96 baseline.

This report provides recommendations for augmenting the existing EOSDIS to handle the Terra mission, improving the partnership between interested parties, adopting a more proactive approach that will better address community satisfaction, and re-examining future data distribution needs.

- The participants recognized that developing the data system for EOS with its unprecedented volumes of earth observation data would be a major challenge for any organization
- It is clear that there is an urgent need for increased hardware capacity in several parts of the system, so that production can maintain a steady and reliable flow of data, even at the Option A+ level of 1X Level 1 and .5X Level 2/3.
- The Option A+ limit on production has insufficient margin for continuous operational production and has prevented the system from establishing critical mass to allow a steady flow of products and runs counter to previous experience on the margins needed for operational production.
- The Insrument teams have made considerable improvements and efficiencies to code and algorithms to minimize the required volumes and loads.
- Based on current knowledge on system and algorithm performance, increased hardware capacity is needed above the 96 baseline to generate the higher-order geophysical data products already commited to for Terra and planned Aqua science. Provisional costs estimates of a three-fold increase in production capacity at the DAACs and SIPS suggest that the required capacities could be obtained by spending approximately 7% of the annual ESDIS budget.
- The impact of forcing the IT's to meet the 1996 baseline capacities appears to be likely to delay scientifically validated data products from Terra by two to three years and may well force similar delays on Aqua data products.
- Reducing volumes and loads by cutting the small number of peer-reviewed geophysical products will prevent EOS from meeting its science goals and is considered as unacceptable by the SWGD.

The primary finding from this working group meeting was that there is a need for a marked increase in the system capacity to generate data products. The current budget situation is difficult on all sides. It is clear that there will need to be frank and open discussions between all of the parties involved in EOSDIS regarding possible options. One area of concern is the impact of the rapid evolution of information technology on the obsolescence of the current system. This impact suggests that we may need to move rapidly from the current system to the more distributed system being envisioned for NewDISS.

1. PURPOSE AND BACKGROUND

It is generally recognized that developing a data system to support processing, archival, and distribution of NASA's Earth Science data is an extraordinarily difficult endeavor. Part of the difficulty lies in the technical problems associated with the very high throughput required of the system, as well as the high processing power, the large storage volume, and the complexity of making the data accessible to a wide variety of data users. Another part of the difficulty lies in the fact that this system must engage a large number of different communities that are unaccustomed to exchanging data with each other. The development process has been difficult and the initial operations have been a learning experience for all parties involved. It is a tribute to the hard work of many members of these communities that EOSDIS is beginning to produce and distribute data.

Shortly before the end of 1999, the EOS Terra platform was launched and EOSDIS began to work with the data from its instruments. These instruments began observing the Earth near the end of February and have continued their observations since. At the March 16-17 meeting of the Science Working Group for the AM Platform (SWAMP), SWAMP members noted that there needs to be on-going evaluation of how well the current EOSDIS can support the Terra, Aqua, and Aura missions for which it has been designed. While there has been considerable discussion concerning the type of Data and Information System (DIS) required beyond the current EOSDIS implementation, there has not been a systematic evaluation of how well the current system is working. Dr. Chris Justice (of the MODIS Land Team) accepted a motion that he should convene a workshop of data representatives from the Terra instrument teams to carry out this evaluation and report to the SWAMP regarding both their concerns and possible solutions.

On June 1-2, the workshop was held at NASA Goddard Space Flight Center (GSFC). It constituted the inaugural meeting of a Science Working Group on Data (SWGD). Participants included (1) Terra instrument team representatives from CERES, MISR, MODIS, MOPITT, (2) EOS Project and ESDIS Project representatives, (3) EOS Project Science Office representatives, including the Terra and Aqua Project Scientists, and (4) DAAC and SIPS representatives. Prior to the workshop, the participants set the following specific objectives:

1) Assess the current status of the data system for Terra, and if possible to assess the planned capacity for Aqua
2) Determine how requirements have changed with respect to the current design and system capacity
3) Identify mismatches between current plans and needed capabilities
4) Recommend solutions to meet evolving needs.

The first day of the workshop was a general session attended by all, with detailed discussion of the DIS status and of the evolving needs of the Terra instrument teams. The second day was a closed session where representatives of the Terra instrument science teams and the EOS Project formulated the preliminary conclusions of this report. The ASTER instrument team did not participate since their relatively modest hardware requirements are adequately addressed by the existing EOSDIS.

2. EVALUATION OF CURRENT EOSDIS STATUS

2.1 Current activities

While the EOSDIS is now operational on its first large-scale application, the Terra mission, its emphasis to date has been on getting the system up and running. This situation will continue for at least several more months. Current Terra data activity centers primarily around product generation, both to bring the system up to operational strength, and to validate the science algorithms. The full range of Terra data products will emerge over the next year or so. If the existing policy of gradually increasing capacity over the next three years is followed to conclusion, full throughput capacity will not be realized until 2003. Distribution of data to Terra science users has not yet been substantially demonstrated, although this will be vital to the long-term success of the system.

2.1.1 Getting the system running

The data system for EOS represents the largest undertaking of its kind to date and continues to be a significant challenge for those designing, building and using the data system. As is typical with a new mission and a new system, the EOSDIS has been subject to a range of problems, which have been compounded by the unprecedented scale of the Terra mission and the complex interdependencies between distributed processing steps. For example for the MODIS instrument there are five distributed groups involved from acquisition to distribution.

Much credit is due to those who have worked and continue to work toward making the system fully operational. They have addressed many of the early operational issues. One particularly troublesome example was the “bit-flip” problem attributed to the on-board solid-state recorder during playback. This problem caused several disruptions to data throughput that substantially reduced early data flow through the system. Problems encountered with the EOS Data Operations System (EDOS) Level Zero Processing System have been due in large part to insufficient hardware and media, as well as inadequate margin for continued production when hardware problems are encountered. With hard work, many of these initial difficulties have been overcome. At the same time, new issues continue to arise and unsolved issues await resolution

Despite the undoubted progress, the system is not yet demonstrating the ability to keep up with daily production even at the current level of 1X for Level 1 products and 0.5X for Level 2 and higher products. Budget cuts and the Option A+ cap on system capacity means that there is no redundancy or backup when hardware is down or new software is being installed. As a result there is rarely complete daily production and there is a backlog of data to be processed. These gaps in production are affecting the early mission science emphasis on product checkout and verification. For example, the large amount (a few hundred orbits) of data for MISR geometric calibration was delayed by over a month because of system issues. The other Terra teams have experienced similar data delays.

Many of the problems to date appear to originate in the EDOS, which includes the Ground System Interface Facility (GSIF) at White Sands and the Level Zero Processing Facility (LZPF) at GSFC. EDOS problems affect not only day-by-day data handling, but also leave a liberal trail of data gaps that hamper production, especially when complete time periods need to be present.

Both CERES and MOPITT base their processing on having 24-hours of data when they start Level 0 to Level 1 production. The EDOS problems have made it difficult for these teams to assemble complete 24-hour periods of Level 0, ephemeris, and attitude data that are used together to produce calibrated and geolocated data within a reasonable (less than several weeks) period of time. Patchiness in daily global production and subsequent processing of data from multiple time periods has slowed MODIS processing. It has also reduced the quality of 8 and 16 day composite products. Product generation at the DAACs and other facilities will be more regular and higher when EDOS production is sutained and consistent. EDOS is systematically undersized, largely as a result of cost cutting. This undersizing also resulted in some Terra data being lost when raw data were not processed to Level 0 before the raw data tapes were recycled, a situation that is unacceptable scientifically and programmatically, a poor advertisement for a mission of this magnitude.

Operational issues within the science processing are more diverse, and many have already been addressed, such as the peculiarities of MISR's orbit-based processing. At the moment, MODIS processing is currently running between two to three weeks behind acquisition due to delay in receipt of Level 0 and definitive ephemeris/attittude from EDOS and Flight Dynamics. However, as EOSDIS operators gain experience, it may be possible to extend our goals. While generation of Level 1 products was removed as a system requirement early in the program, it is possible that additonal operational experience would make it possible for EOSDIS to become responsive to timely events such as monitoring the fires in the Western US. If such a goal could be achieved, it would allow EOS to respond to disasters and natural hazards applications.

Funds have been made available for additional hardware to relieve some acute capacity issues, such as those at the GSIF and for MODIS processing at the GSFC DAAC. However, the workshop participants maintain that additional hardware capacity is essential for sutained operational production as we explain below.

2.1.2 Summary of Production Status and Issues by AM Instrument

2.1.2.1. ASTER

Data transfers from the U.S. to Japan and back have been a matter of some concern.

2.1.2.2. CERES

CERES processing needs three sources of data. The table below shows the percentage of missing data for each as of 8/23:

% Missing	Feb	Mar	Apr	May	Jun	Jul	Aug
L0 (FM1 & FM2)	4	0	1	0	13	39	26
Ephemeris	0	0	0	2	10	24	1
Attitude	0	1	1	4	9	26	7

Due to one or more missing source, CERES can only process 58% of their data for the entire month of July. However, in the last couple of weeks, all sources have been timely and complete.

2.1.2.3. MISR

The projected shortfall for MISR processing is primarily in the area of processing capability. At present, the hardware architecture can sustain about 65% of Level 1 processing and 10% level 2 processing. This system was originally sized for 100% level 1 and 25% level 2. Currently, decisions are made as to how to fit within the current processing envelope.

Some of this processing shortfall is expected to be alleviated when the planned upgrade to an Origin 3000 occurs. At that time, the system is expected to be able to support 1X processing equivalent to 100% of Level 1, plus 50% of Level 2, plus an as yet undetermined amount of Level 3. The sizing of this system does not compensate for level 0 data not being shipped to the DAAC in a timely manner and maintenance time required at the DAAC. When data does not arrive at the DAAC in a timely manner, a backlog occurs. As the backlogs increase, decisions must be made as to what data does not get processed due to processing constraints.

2.1.2.4. MODIS

The MODIS production system has a large number of distributed and dependent components, since five to six different groups participate in the data stream. All of these components need to be working reliably and consistently to provide a regular flow of data. There is no redundancy in the hardware, so that when one piece of the system is down, production is commonly interrupted. Sharing of processing resources between the MODAPS (L2/3) and the GDAAC (Level 1), and tapes between MODAPS and EDOS has provided temporary reduction of bottlenecks in Level 1 production. MODIS production is currently running approximately three weeks behind instrument acquisition. There have been a large number of temporal and spatial gaps in the global daily coverage which is has hindered production and evaluation of the time-series products. Product validation has been held up by limited data availability. Delays in data flowing from EDOS to the GDAAC has meant that multiple periods are often being processed at any one time. This, in turn, places a large burden on the online and near-line archive and data retrieval. Improvements are being made to the algorithms based on instrument performance and new code is being promoted into the operational chain. The capacity for reprocessing even small amounts of data from selected validation periods is extremely limited. Reprocessing to generate time-series data using revised algorithms will have to wait until the necessary hardware is obtained. MODIS Level 1 data were made available to the general public by the GSFC DAAC in May. Preliminary Level 2/3 land products are currently available through the EDC DAAC and the ESIPS supported 250m system as part of a phase-in of production and distribution.

2.1.2.5. MOPITT

The early behavior of the system produced a larger than expected number of intermediate files, but no major problems appear to have affected MOPITT data production.

3 EVOLUTION OF REQUIREMENTS

The original EOSDIS concept has evolved considerably since its inception. This was inevitable for meeting the evolving needs of the instruments and is accelerating. It is most apparent where instruments have elected not to use the EOSDIS Core System (ECS) for their processing, but instead to build and operate their own Science Information Processing Systems (SIPS). Examples are the MODIS Data Analysis Processing System (MODAPS), which currently generates MODIS Level 2 and higher-level products, and the Langley Terra Information System (LaTIS), which handles processing, archiving, and distribution for CERES. Even where the original ECS concept remains as intended, system internals have evolved significantly. Local extensions have addressed instrument-unique capabilities not in or deleted from the original concept, such as the special quality assessment (QA) database for MISR.

ESDIS has funded and supported these as they have provided affordable and desirable solutions to known data processing challenges and have been an extremely effective and popular way of harnessing expertise within the science community to ESDIS challenges. According to the ESDIS presentation on Option A+, this "adaptive processing" approach saved $22.3 million dollars over the cost of completing development and then operating the ECS to produce those products currently made in SIPS. In most cases the SIPS have been closer to the instrument teams which enables science decisions to guide prioritization of effort. In the case of MODIS, moving to a SIPS resulted in a 4X speed-up in the integration of science software changes. It also provided the capability to distribute products on DLT and AIT-2 tapes and enabled enhancements such as selective distribution based on spatial coordinates, Q/A metadata, as well as support for subsampling and subsetting products.

Continued evolution of the data system is critical as requirements unfold, as the initial implementation becomes obsolete and uneconomical, and as new information technology becomes available. The SIPS and DAACs need the flexibility to enhance the trade-off of needs and resources. Both the DAACs and the instrument teams have demonstrated substantial initiatives in exploring new processing systems and systems architectures. These innovations include Goddard's Simple Scalable Script-based Science Processing (S4P) and the MODIS Linux clusters. The SWGD is very encouraged by these explorations and feels that the DAACs and SIPS should be provided an even greater degree of autonomy and independence in the future. In addition, the NASA ESIPS program has proven a useful proving ground to providing new approaches and technology for data services. Directing the ESIPS to address some of the data challenges facing EOS may provide some of the additional innovation that is needed, particularly with respect to distribution and data services.

3.1 Under-scoping of the Terra system

The current EOSDIS has been implemented in a fashion that does not satisfactorily account for the inevitably evolving needs described above. The system is sized to a baseline frozen in February 1996, based on figures derived by the Ad Hoc Working Group on Production (AHWGP). The numbers in this baseline date from 1995. While they were a best effort at that time, many of the algorithms on which they are based were not then well understood, nor was the system environment well known or stable. As the system has developed, factors that determine the total throughput have become clearer, such as system overhead, the workforce needed to plan jobs, and the nature of data staging and the sequencing of production. These factors were neither known nor included in the baseline estimates..

Both the science algorithms and the system issues are now much clearer, since the IT's have a moderate amount of operational experience. This knowledge provides a timely basis for revising the system. At the same time, there are areas yet to be implemented where the needs are not clear. In particular the distribution system has yet to be fully exercised. These needs will have to be addressed later.

Since ESDIS established the February 1996 baseline, the production and storage capacity of EOSDIS has not been updated based on formal consultation with the science teams and the evolving clarification of algorithm resource use. On the contrary, the system requirements since that time have been made primarily in response to budget issues. The base size of the system is now derived largely from the so-called "Option A+", which includes only a 1X distribution of data to users, and only a 2X reprocessing capability once full capacity is reached in 2003. This model does not reflect NASA experience in past missions. For example SeaWiFs runs at over 100X production and TRMM at 10X. The instrument teams were not consulted on the likely impact of Option A+ on product generation. Concerns that were raised at the time by MODIS to the Reber/Barron Committee went unresolved Now that we are facing the results of implementing Option A+, these issues are returning to the fore. At the last IWG, the lack of global production of MODIS land products was identified as a major concern. Nonetheless, the hardware capacities continue to be based on the baseline of February 1996, even though it is clearly no longer scientifically appropriate.

3.2 Need for increased production capacity

At the workshop, Terra instrument team representatives presented revised system capacity estimates for their algorithms, thereby providing a comparison with the February 1996 baseline. These figures and the associated explanation are given in Appendix C. In summary they indicate the following requirements:

- MODIS, CERES, and MISR computational needs are larger than the baseline by a factor of about 3.
- MODIS archiving volume for Level 2 and above are larger than the baseline by a factor of about 3.
- MODIS ingest needs are about 5 times larger than the baseline.
- CERES ingest needs are about 3 times that of the baseline.
- MISR storage needs have decreased due to data compression working more effectively than the baseline assumptions.
- ASTER has no significant change in requirements.

These changes are neither surprising nor excessive given the level of understanding of the interaction between the algorithms and the production environment that was not well defined, let alone operating four years ago. Indeed, it appears that the estimated hardware capacity to produce a month of CERES-TRMM data has dropped by about a factor of two and one-half over the time period from 1995 to the present (see fig. C1 on p. 20).

The constraints of the Option A+ model remain to be dealt with, although this model defines the formally approved system capacity. It cannot satisfactorily support the implemented algorithms. It should be appreciated that the science software developed by the instrument teams is extensive and represents a major investment. In more than one instance, the code size is comparable to that of the ECS. The instrument teams have already reduced the hardware requirements of their algorithms and have streamlined their code to reduce the processing volume. They have also organized a phased approach to product generation and distribution to alleviate pressure in the early days of production.

To remain with Option A+ is nearly certain to delay production of most of the validated scientific data products by two to three years. This option may also imply a need to either process less than the full Terra data set, or to further simplify Terra algorithms (at the cost of increasing the errors in the products) and to reduce the scope of the user communities. Clearly, processing less than the Terra data set does not meet EOS commitments. Increasing error in the data products and reducing the scope of the user community will also adversely affect the reputation of EOS. Adoption of increased processing capacity is therefore indicated. It is clear that the cost savings attained through squeezing the hardware budget have imperiled the scientific scope of the Terra mission. Hardware has always been a small part of the EOSDIS budget (c. 15%) . To remove flexibility in this area will make it impossible to achieve the mission science goals.

4 NEAR-TERM RECOMMENDATIONS

This section proposes specific actions for the immediate future to meet the current commitments of the EOS missions to their respective end users. Recommendations concerning EDOS and the distribution system are not included here, although its sustained operational production and easy access to data are central to overall system success. The SWGD assumes that ESDIS and its contractors are working on solutions for EDOS, and that data distribution will be a topic for future attention by the working group.

4.1 Update system capacities to reflect improved understanding

The SWGD will work with the DAACs, SIPS, and ESDIS to determine the new hardware capacity needed to update the capacity estimates prepared for the February 1996 baseline. The figures in Appendix C show preliminary estimates of the new requirements for CERES, MISR, and MODIS. Spreadsheets with more detailed quantification of requirements have been provided to ESDIS and several of the DAACs. The SWGD expects to revise these estimates at least semi-annually.

The initial rough estimate of the added requirements for Terra and some of Aqua appears to be equivalent to about seven percent of the current annual EOSDIS budget or about $14 million per year for the next two years. In return, the system capacities would be expanded by a factor of three. The SWGD has received data to confirm these estimates from ESDIS and, independently, from the DAACs. There are several issues that need further clarification, although the cost of computers, disks, and tape silos to meet the need identified in Appendix C appears to be reasonably well understood. Keeping these resource estimates up to date will be an on-going activity for several reasons:

- The increased hardware capacity would be built up over a few years and probably include replacement of now-obsolete hardware.
- It needs to address production needs not only for Terra, but also for Aqua, Aura, and other EOS missions.
- It needs to address reprocessing adequately and in particular the time it takes to reprocess a time series record.
- It needs to address evolving needs during the missions.

The resources required to implement the updated baseline require assessment by the instrument teams, DAACs, ESDIS, and other relevant parties. Because hardware sizing does not generally involve system restructuring or redesign, because the required resources appear modest compared to overall system costs, and because this issue warrants special attention, it may be appropriate to see if the required resources can be arranged separately from other EOSDIS resources. In this time of budget limitations, cost effective ways of providing the necessary hardware should be pursued. The instrument teams and the DAACs both need to be party to any reprioritization of existing funds if that is necessary to obtain the augmentation.

This augmentation also must include a reassessment of reprocessing in order to overcome the substantial limitations of Option A+. Gaps in the time-series record since launch and major improvements to algorithms based on actual instrument performance make reprocessing essential to meet the science goals for the monitoring instruments e.g. MODIS and CERES. Part of the difficulty we have encountered with Option A+ may arise from overly simplistic budgetary models of data production and reprocessing. The SWGD will explore more rigorous approaches to tying data production and validation schedules to capacity estimation. As it gains experience with these approaches, the SWGD may be able to provide improved estimation methods to their colleagues on other teams. In addition, the instrument teams should look at new approaches to reprocessing that may improve product quality or system efficiency with reduced system capacity.

The scope of the increased capacity to produce, archive and distribute data should embrace Terra, Aqua, and other instruments in the EOS complement. The SWGD notes that the needs and priorities of individual teams differ markedly. Given the success of the local adaptations the teams have made to their situations, more localized approaches appear appropriate to avoid the difficulties of the “one size fits all” approach of the original system concept. It is appropriate to involve the emerging DAAC alliance that is addressing future needs of the system. Local revision of the system must be possible throughout the mission, because needs will evolve locally.

We also note that hardware updates may force basic changes to the processing system. These changes need to be developed as appropriate. They may move beyond the current ECS paradigm where economic and practicality issues so indicate. Such changes might include migration to new systems or evolution of existing ones. The need for them will vary from instrument team to instrument team.

4.2 Begin a new partnering approach between all interested parties

There is a vital interdependence between the ESDIS Project and the EOS scientific community, especially the EOS instrument teams. Too frequently there has been insufficient communication between them on decisions that impact both sides. The degree of consultation between ESDIS and the instrument teams regarding decisions and plans needs to improve. An example would be more open discussion of the trade-offs involved in resource allocations, including the determination of hardware capacities within the overall EOS budget context.

The SWGD provides a mechanism for this. However, it should function as an independent body and may be rescoped from time to time to suit the needs of current missions. It will work in concert with other bodies such as the DAAC alliance and the DAAC user working groups

- to provide improved science guidance to ESDIS,
- to provide an assessment of current EOSDIS operations to the ESSAAC subcommittee on data issues, currently chaired by Sara Graves, and
- to provide a perspective on the instrument team operations to NASA Headquarters.

The SWGD does not replace existing long-term planning bodies such as those developing the concepts of NEWDISS. It will function satisfactorily only if it is constituent driven. In the context of EOS data it will be necessary for the SWAMP Working Group on Data to expand to include the instruments from the other EOS platforms. Within this enriched context, and indeed to enable it, the instrument teams, in concert with the DAACs, must be allowed to accept increased responsibility and accountability. They have already demonstrated a reasonable level of success in dealing with both production and distribution issues. Increasing their autonomy and independence will be critical to responding to future challenges.

4.3 Adopt a proactive EOS approach to providing community satisfaction

The EOS mission is predicated on expanding the availability of validated scientific data to a much wider community than heretofore. In order to achieve this goal, the full community that interacts with the mission, including data users, data producers, and NASA management, must feel that their needs are being met. The SWGD feels that the satisfaction of this community requires a more proactive approach than the one we currently have. There are many ways we can improve in this area, including the following:

Develop ways of more rapidly evolving critical systems and services even while working within increased budget pressures. There is already evidence that this can be done. For example, MODAPS processing capacity was recently augmented by flexible cost arrangements between ESDIS and the MODIS team. The MODIS 250m Land production and distribution system has been developed in the context of the ESIPS program. Likewise, the full (100%) Option A+ complement of hardware for MISR processing was recently put in place at the LaRC DAAC. However, there are areas that have not responded to the evolution of the community needs. For example, over the last several years the ECS has remained dedicated to 8 mm tape for product distribution, even though the user community has long since moved to higher-capacity and faster media. Such areas need attention.

Encourage development of new production architectures. There will be instances where this approach is the only practicable path to enhancing performance while reducing costs and development time.

Encourage development of new, innovative archive and distribution mechanisms. While data distribution limitations have not yet emerged, there is great concern that they will. The next section of the report makes more detailed suggestions in this area.

Implementing these recommendations requires attention to the needs of individual data producers and data users, rather than a blanket approach to all. Nevertheless, the different parts of the EOS community have much in common. The SWGD expects to meet with the DAACs, SIPSs, and ESDIS on a semi-annual basis to identify new opportunities and technologies for infusion into EOSDIS. Perhaps meeting outside the D.C. area would encourage dialog with the wider user community.

4.4 Examine future distribution needs and capacities

While the current limits to distribution have caused no major user complaints so far, there has been very limited data ordering beyond that by the current scientific teams conducting data validation. The pre-Terra experience in the relative fraction of data delivered by media and over the net varies from DAAC to DAAC, ranging from 50% media / 50% net at GSFC to 90% media / 10% net at LaRC. In addition, experience suggests that a total distribution in the vicinity of 5X will be needed to meet realistic expectations of user needs.

There are innovative approaches to meeting data distribution needs that may be cost-effective. For example:

- Ftp caches of scientific validation data sets and WWW caches of images for public distribution. Moving these data directly from staging disk to the caches as part of planned production would ease the burden on the secondary and tertiary storage devices (tape silos) in the main archives.
- Simple web pages for searches by knowledgeable data users in conjunction with ftp caches. This approach should be relatively inexpensive and not affect the fundamental searchability of the EOSDIS data.
- Subsetting capabilities appropriate to each user community. Such subsetting offers substantial possibilities for reducing distribution load from the data centers.
- "Guest investigator" access for high volume users, where scientists can visit DAACs or SIPS and obtain data through direct connections to LANs within the data centers.
- "Lending libraries" of data on media that do not need expensive archive access or network transfers.

The SWGD recommends that the DAACs, their user working groups, and ESDIS work together to evaluate these suggestions (and others that may appear) with regard to practicality and cost, hopefully moving toward rapid implementation of those that are most promising. Initiatives of this type are already evident at the GSFC and LaRC DAACs and could provide a basis for the evolution of EOS into NASA's NEWDISS program.

The SWGD also recommends that there be a meeting about six months from now to address data distribution status and archive access needs. This could take the form of another SWGD workshop. The workshop should cover both quantitative data on current distribution and rigorous user modeling of future needs. It should include all interested groups, especially the DAACs and their user groups.

5 FUTURE CHALLENGES AND OPPORTUNITIES

The recommendations above cover the most immediate issues identified by the workshop. They are part of a greater range of challenges and opportunities for improvements to the service EOSDIS can provide to its user community. Some specific challenges for the coming months and years that were discussed at the workshop include:

- Reprocessing
- Perturbations to operations, particularly those involving the Aqua turn-on and the associated increase in volumes and loads.
- Data Distribution
- Dealing with the Diversity of the Producer and User Communities
- Obsolescence
- Transition to a non-ECS Environment

Some of these have been addressed wholly or partly in this report. All provide scope for on-going consideration by the SWGD and others to foster and maintain an EOS data system that appropriately satisfies the user community needs.

6 CONCLUSIONS

This workshop provided a useful forum for sharing views and experiences among the DAACs, ESDIS, and the instrument teams. It succeeded in identifying more accurate and meaningful system capacity requirements than are currently implied by the existing baseline, which dates from February 1996. The SWGD expects to help maintain these estimates in the future in conjunction with ESDIS, the DAACs and the SIPS. There is a need to extend the SWGD type of activities to include the Aqua and Aura teams. Neither of these teams was able to participate in the workshop.

The workshop was able to provide specific recommendations:

- for addressing near-term issues beyond the baseline capacities
- for improving the partnership between interested groups such as ESDIS, the instrument teams, the DAACs, and the user working groups, and
- for a more proactive mechanism to address community satisfaction

After the DIS settles into more a more stable operational condition, the SWGD plans to re-examine future EOS data distribution needs and capabilities in conjunction with the other interested parties.

APPENDIX A: WORKSHOP ATTENDEES

- Bruce Barkstrom, CERES
- Ed Masuoka, MODIS
- Mike Moore, ESDIS
- Vanessa Griffin, ESDIS
- Steve Fox, ECS
- Claire Parkinson, Aqua
- Jim Galasso, LDAAC
- Graham Bothwell, MISR
- Robert Wolfe, MODIS
- Earl Hansen, MISR
- Richard McGinnis, LDAAC
- Steve Kempler, GDAAC
- Skip Reber, ESDIS
- Jim Ranson, Terra
- Chris Justice, MODIS
- Francesco Bordi, EOS
- Douglas Jaton, EDAAC
- Hampapuram Ramapriyan, ESDIS
- Dan Ziskin, MOPITT

APPENDIX B: WORKSHOP AGENDA

Pre meeting:

- Submission of requested material on volumes and loads

One day open meeting:

- Compiling improved information
- Developing a revised database on needs versus current and planned capacity - devise common reporting /display of material
- Identifying any major mismatches
- Identifying current bottlenecks
  - Suggested solutions to current bottlenecks (near-term)
- Identifying anticipated challenges and improvements
  - Suggested approaches to challenges
- Suggested options for meeting our goals
  - What additional resources are needed for hardware capacity increases?
  - Are there cost saving approaches and solutions that could be adopted?
- Identify next steps – additional issues that need to be addressed

Half day closed door session for instrument teams:

- To draft a formal letter and recommendations from the sub group meeting to AM Project Scientist to send to the Project

APPENDIX C: CHARTS SHOWING DRAFT JUNE 2000 BASELINE

CERES Requirements

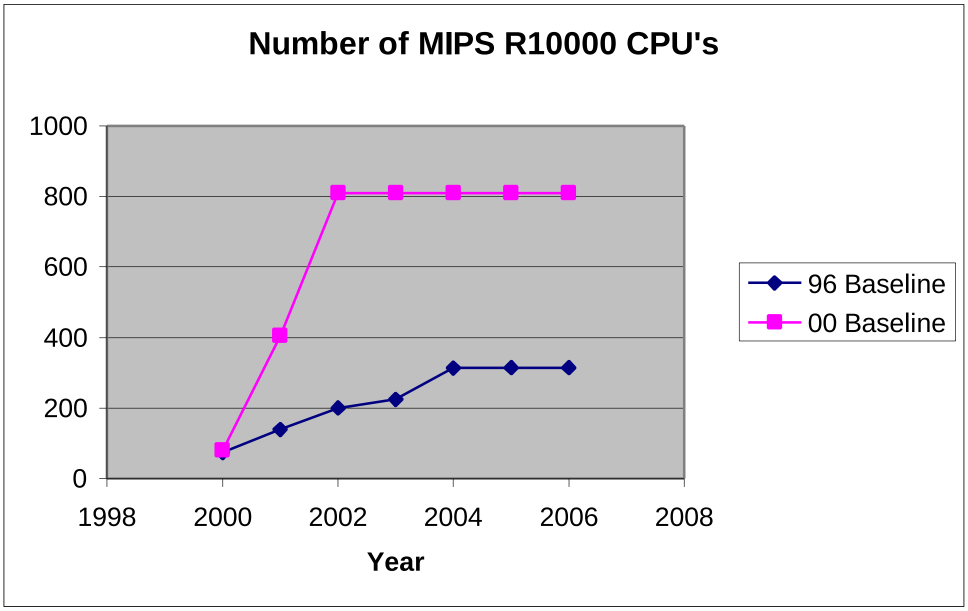

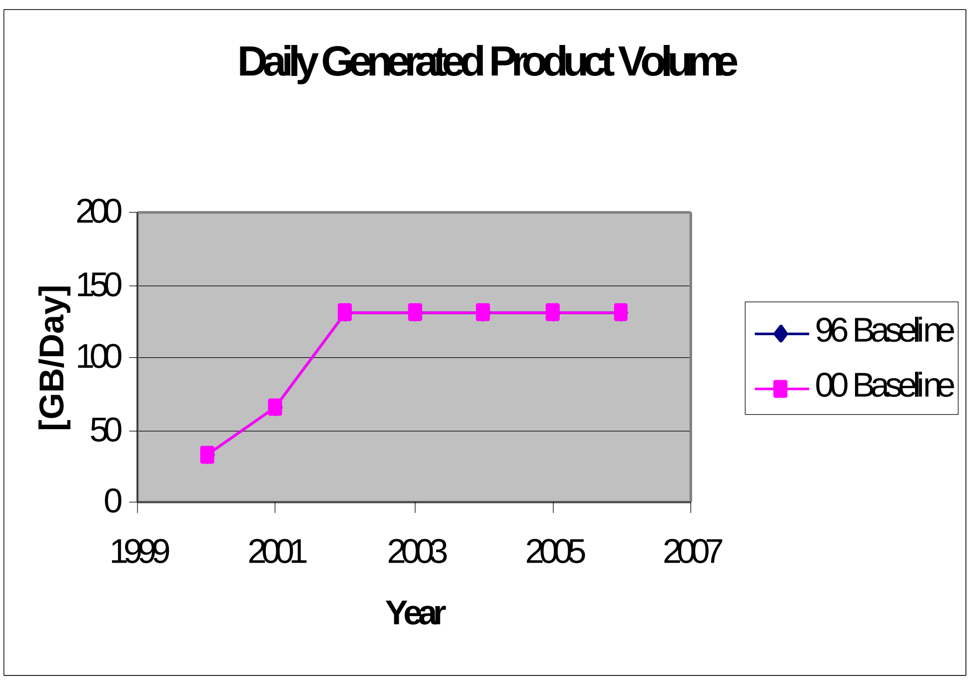

MISR Requirements

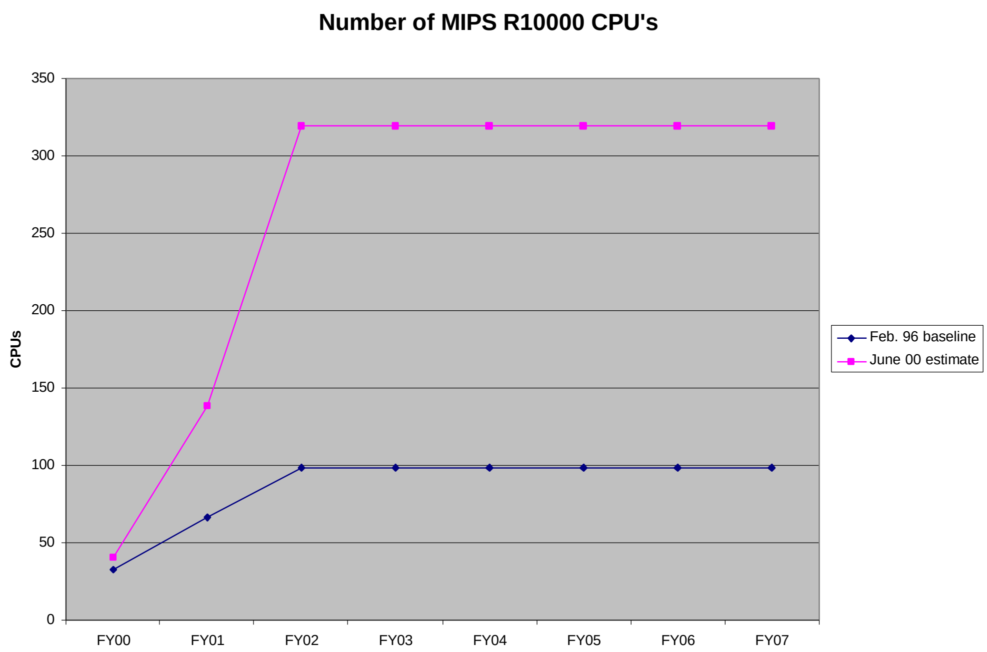

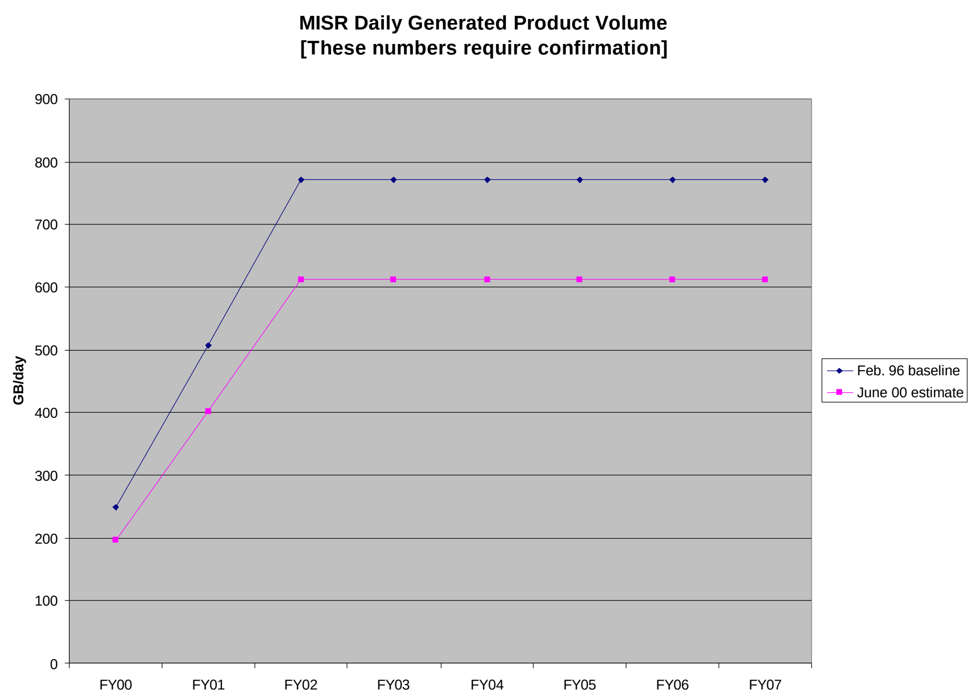

MODIS Requirements

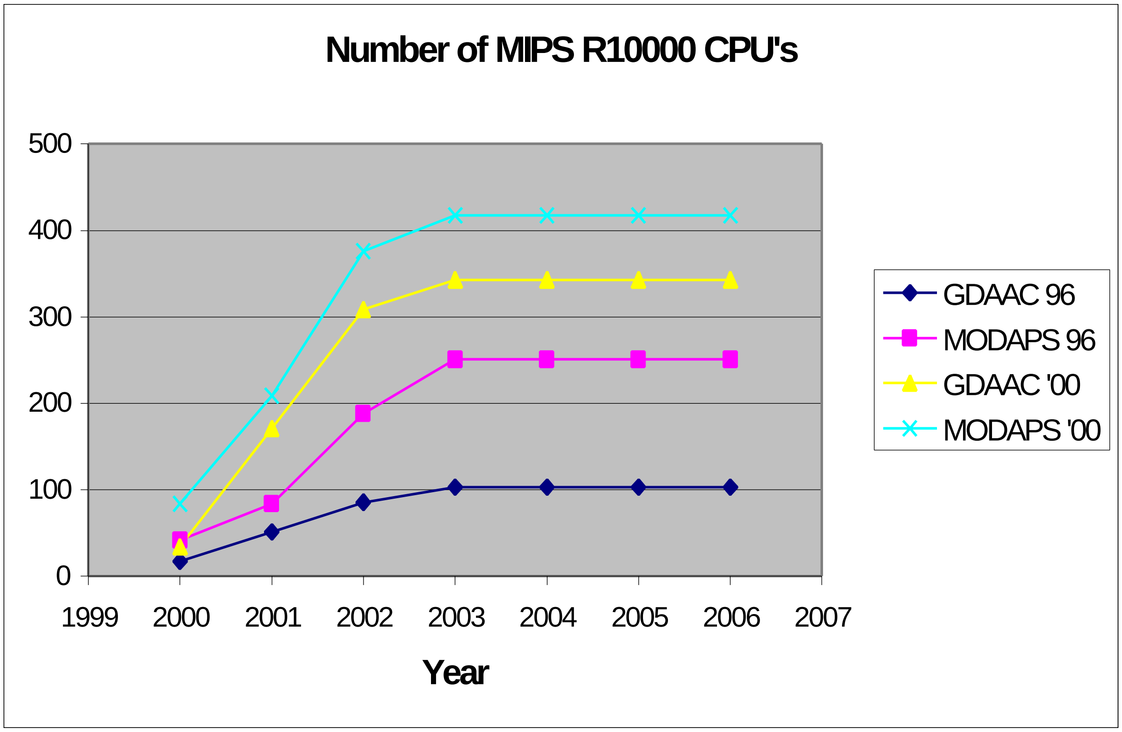

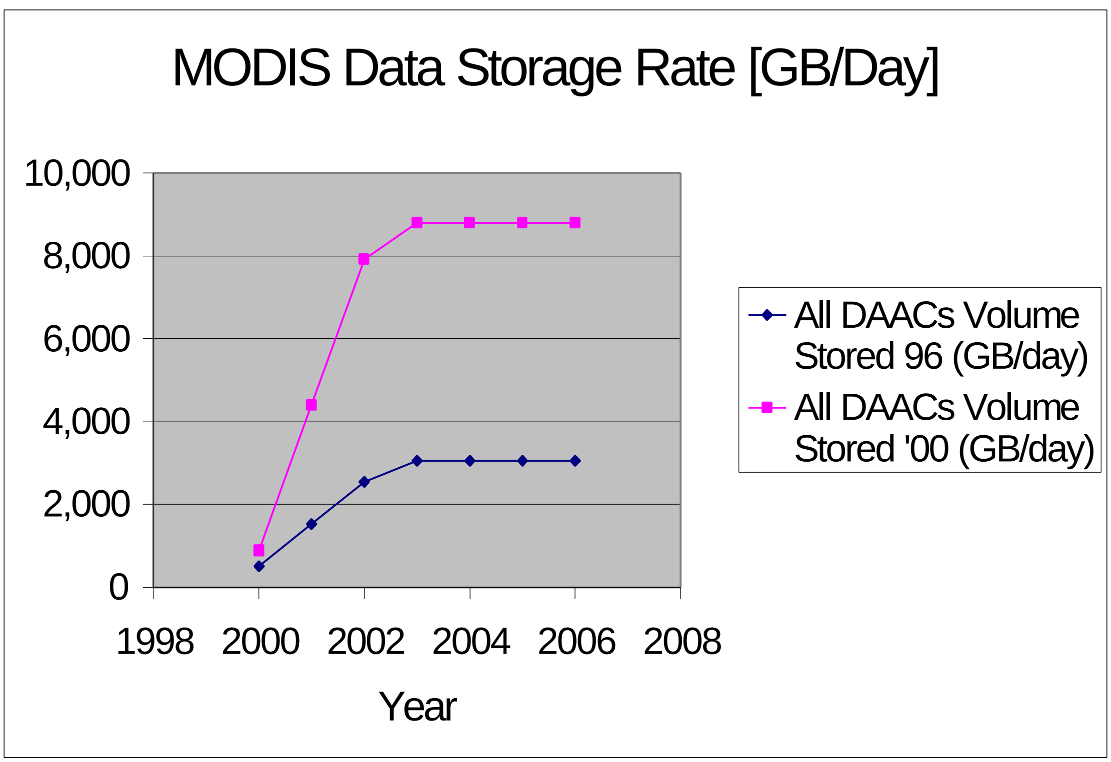

CERES Science Impacts

CERES provides a peer-reviewed science investigation in five areas:

- Continue ERBE measurements of Earth radiation budget and cloud forcing
- [*] Improve scene identification of radiation budget with simultaneous and collocated imager data (VIRS and MODIS)
- [*] Provide new Angular Distribution Models to cut instantaneous ERB flux errors by factors of two to four
- [*] Provide empirical surface radiation budget fluxes
- [*] Provide vertical flux profiles for shortwave and longwave radiation based on observationally derived cloud properties

The bullets marked with the [*] are new areas of Earth science research that are intended to significantly advance the understanding of the interaction between clouds and the Earth's radiation budget. This interaction remains the largest source of uncertainty in our understanding of how the climate will respond to various kinds of perturbations.

The CERES measurements require simultaneous and collocated cloud imager data because clouds and water vapor change significantly over time scales of minutes. On Terra and Aqua, CERES has two instruments – one devoted to measuring the spatial distribution of radiation fluxes and the other devoted to measuring the full angular distribution of outgoing radiances. A single CERES instrument also flies on the TRMM spacecraft. This combination of instruments on different spacecraft is necessary to reduce the temporal sampling errors that affect the monthly averages that are one of the primary data products from the CERES investigation.
In addition to the large investment in satellite instruments, CERES has also developed a significant investment in the source code for its data production efforts. There are about 600,000 lines of code for the algorithms and about 45,000 lines of production scripts to control data production. This investment in software is not unusual for the EOS instrument teams.

Because CERES has been reducing data from the TRMM instrument since it was launched in late 1997, CERES also has a considerable investment in data validation. This investment is expanded in the Terra and Aqua eras by using data from the Department of Energy's Atmospheric Radiation Measurement (ARM) program in the Southern Great Plains, Kwajalin, and North Slope Alaska sites – as well as the Baseline Surface Radiation Network.

Why has there been an increase in CERES CPU capacity needs and what has CERES done to reduce these needs?

The answer to that question begins with the process by which CERES has developed its capacity estimates. A significant part of the CERES investigation arises from its history of these measurements in the Earth Radiation Budget Experiment (ERBE). Two of the thirteen CERES subsystems reuse code from ERBE and provided good estimates of the capacity needed for this work. In the early years of EOS, the CERES team estimated CPU power by scaling from several similar algorithms, using number of pixels as the fundamental sizing parameter.

As time has gone by, the CERES Data Management Team has kept careful records of how the actual processing capacity matches up with the estimated capacity. Figure C1 demonstrates that the CERES CPU requirements on TRMM have actually gone down as the team has gained experience – contrary to the nearly universal expectation that instrument teams are gluttons for processing power.

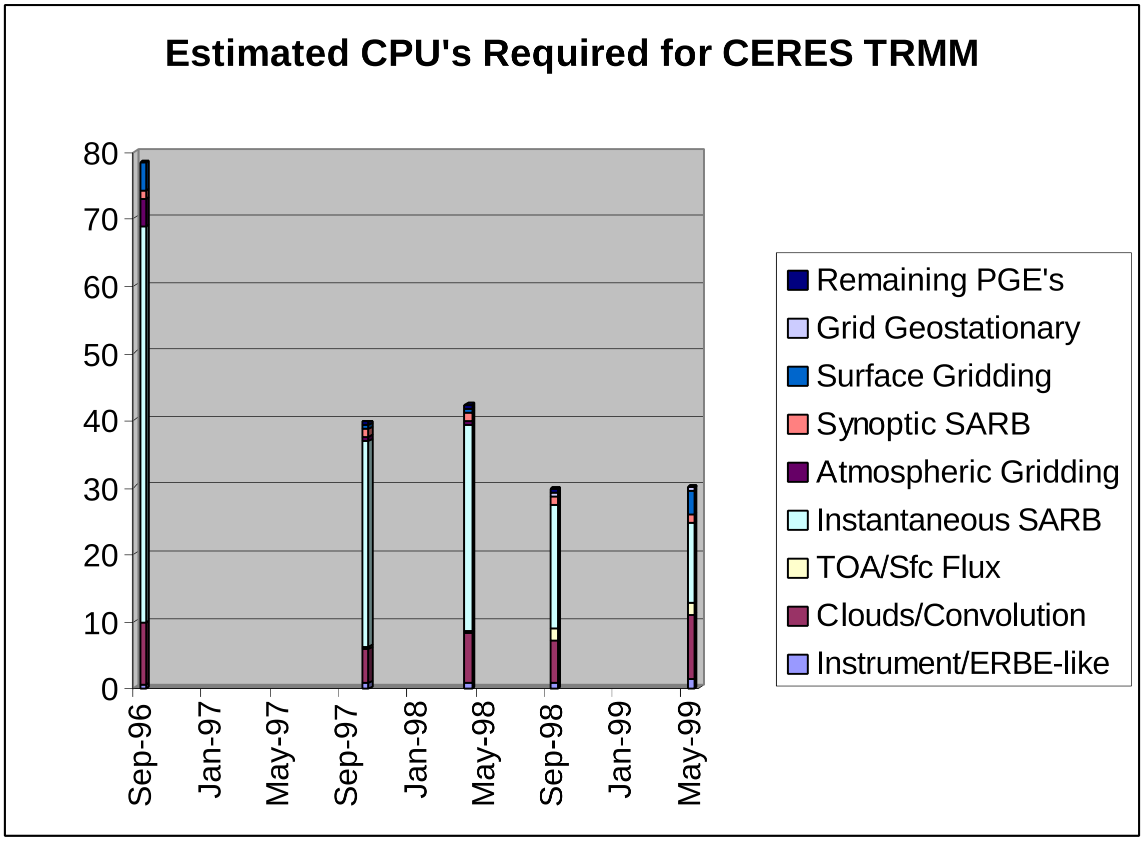

Figure C1. Estimated Number of MIPS R10000 Chips (185 MHz) to Process One Month of all CERES Data Products on TRMM.

The two columns on the left of this figure are pre-launch estimates obtained during the two software releases the CERES Team produced during that period. The three columns on the right provide the measured CPU needs based on the optimized run-time experience in the LaTIS production environment. The most expensive subsystem (Instantaneous SARB) was identified early and targeted for capacity reduction before launch.

In additon to this experience, CERES has had a number of successful capacitty reduction exercises:

- Ingest - **down 6X**
  - Original ESDIS/ECS ingest estimates included 1 MODIS stream for each CERES instrument
  - CERES recommended not only reducing to 1 MODIS stream per satellite but also subsetting to only 1/3 of channels
- SARB Algorithms - **down 10X**
  - Most CPU-intensive algorithms for TRMM
  - CPU needs reduced by factor of 10

- Personal efforts by CERES PI on AHWGP - **19X reduction in ESDIS/ECS working groups**
- Cloud Algorithms - **work in progress, but scientific risk**
  - Considering reduced imager resolution (may increase spatial noise)
  - Working on possible removal of clear-sky reflectance history (may reduce cloud detectability)

Based on the fact that the CERES capacity estimates have been restrained over time and that the CERES team has acted responsively to bring these capacity requirements under control, it seems reasonable to identify the cause of the apparent change in CPU requirement as the change in the project capacity profile, not in the CERES needs.

What would be the science impact of not getting additional CPUs?

It is clear that if CERES cannot obtain increased CPU capacity, the schedule for production of new scientific information will slip, probably by several years. In addition, we can identify the following consequences if the CERES data products were drastically reduced:

- Critically reduces ability of the EOS program to reduce the uncertainty in cloud-radiation interaction in a timely fashion
- Effectively reduces ability of EOS
  - to provide improved long-term cloud properties
  - to reduce uncertainty in angular distribution models and TOA fluxes
  - to improve surface radiation budget
  - to provide new information on atmospheric energy budget
  - to validate cloud/climate model radiative fluxes segregated by cloud type and cloud properties
- Probable reduction in EOSDIS user community
- Slows or reduces improvements in commercial data products used by solar energy and home building industries with international ties
- Reduces ability to provide timely data to students who participate in Student Cloud Observations On-Line (S'COOL)

MISR Science Impacts

Meeting MISR's mandated objectives requires an upgrade in processing power by about a factor of three. One of the principal impacts of the projected shortfall is in reprocessing capability. The ability to reprocess data is essential for developing effective MISR algorithms. No data with the range of view angles, precise radiometric calibration, and relatively high spatial resolution produced by MISR has ever before been obtained from space. Entirely new retrieval approaches have been developed to take advantage of the instrument's global capabilities. The result is a range of preliminary new atmospheric and surface products, providing global measures of tropospheric cloud heights and winds, cloud albedos classified by cloud type, aerosol optical depths classified by particle type, and surface biophysical parameters based upon physical models.

These retrieval algorithms must be tested and refined using actual data, an effort that began as soon as MISR data became available. Most or all of the MISR data needs to be processed through Level 1B2 to identify and obtain high-quality coverage for the range of natural conditions to be studied. New Level 2 algorithms must be run and re-run as the results are analyzed, and as modifications to the algorithms are made. Based on experience with previous satellite remote sensing experiments, this process will continue for at least a few years. If CPU throughput becomes a factor adversely affecting our schedule, the delay will make it more difficult for the MISR Team to complete its work, and it will postpone, and may compromise, some key scientific benefits of the MISR mission.

MISR has a mandate to generate global products at Level 2 and Level 3. Due to the limited swath width of MISR, the coverage allows for statistically meaningful measures of climate change variables on monthly temporal scales, but only if the entire MISR data stream is processed. These are the shortest scales over which global climate change can reasonably be assessed. Ongoing algorithm development, refinement, and validation requires that we have the throughput to reprocess several times. Given these realities, taking three years to process a single year of data is unacceptable for timely inference of Terra global climate results, and it adversely impacts the ability to use such knowledge to plan for future Earth science mission needs.

Less than full production of Level 2 data also compromises the experiment by requiring a labor-intensive and continual replanning of processing priorities in order to catch target of opportunity such as volcanic eruptions, fire outbreaks, or dust storms. This process can incur processing delays that limit the timeliness of the products for use by the scientific community as well as other users with rapid turn-around demands, such as the media. We have already experienced requests for products with less than 1 day's notice, and the geographic location of such requests cannot be predicted in advance.

Resiliency in the MISR data processing system is required to deal with partial system outages. Additionally, innovative uses and approaches to processing MISR data are likely given the newness of the experiment and richness of the data content. It is unrealistic to expect that the current set of algorithms and products represents the entire application of NASA's investment. Inputs from the science community, external to the MISR Science Team, could drive the processing requirements in as yet unknown ways. Therefore, the data processing system must also contain sufficient resiliency to accommodate new needs.

MODIS Science Impacts

Why there has been an increase in MODIS product volume and processing loads and what would be the science impact of not getting additional archive volume?

The archived daily product volume per MODIS instrument in the February 1996 baseline is 537GB per day. It is been divided between MODIS science disciplines as follows:

- 312GB for Level 0 and Level 1 products (includes Cloud Mask)
- 29GB for Level 2&3 Atmosphere products,
- 26GB for Level 2&3 Ocean products,
- 154GB for Level 2&3 Land products exclusive of snow and ice

- 16GB for Level 2&3 Land snow and ice products.

Current average daily product volumes per MODIS instrument is 822GB per day split between the science disciplines as follows:

- 295GB for Level 0 and Level 1 products (includes Cloud Mask),
- 18GB for Level 2&3 Atmosphere products,
- 188GB for Level 2&3 Ocean products,
- 281GB for Level 2&3 Land products exclusive of snow and ice
- 40GB for Level 2&3 Land snow and ice products.

Growth in archive volume for MODIS can be explained in part by growth in the size of individual products through the expansion of per pixel science quality assurance metadata fields, inclusion of useful new parameters within products and the SWAMP mandidate that all Terra instruments generate these fields. In general, the growth of individual products has resulted in only modest size increases. The major reasons for increased requirements for archive storage between 2/96 and today are science discipline specific and will be explained below for ocean and land products.

MODIS Ocean Product Volume Increases

MODIS ocean product volume is 29GB per day in the February 1996 baseline. Prior to the AHWGP meeting at which the baseline was established the MODIS Ocean Science Team estimated the volume of MODIS products to be 110GB per day. Four years ago we estimated that 110GB per day (give or take 10%) was sufficent to store MODIS standard ocean science products. Though MODIS products weren't being produced at that time, the experience of MODIS Ocean Team members with making global oceans products led to a good estimate of the archive storage required for these products. However, three changes occurred which resulted in our product volume, 188GB per day, being signficantly larger than the 29GB per day in the 2/96 baseline.

First, the archive storage for Level 2 Ocean products was not included in the 2/96 baseline. Dr. Robert Evans had agreed that the Level 2 products could be dropped from the archive if the EOSDIS system would take on the task of producing any Level 2 Ocean products requested by scientists and shipping them to the requester within a reasonable time. The archive storage was dropped at that time but the capacity to do on-demand processing of requests for days or weeks of global data was never implemeted in the ECS nor included in the resources for the MODAPS SIPS.

Second, MODIS Ocean products were originally implemented as collections of parameters which shared a large block of common data. When ECS developers were unable to implement the capability to ship individual parameters to end-users from a single file which was a collection of parameters, they requested that the MODIS ocean team split each parameter out as a separate file. This led to a sizable increase in the number of files and a very significant increase in archive volume as the common elements in the original multi-parameter file had to be duplicated in each of the new single parameter files.

Third, at a Science Working Group for the AM Platform (SWAMP) meeting, the EOS-AM Project Scientist levied the requirement on all instruments to produce global products at common resolutions on an equal angle grid (1 degree, half degree or quarter degree Climate Modelling Grid) products. Implementing these products did not lead to a large increase in volume but did add hundreds of additional files per day for ingest at the GDAAC.

Fourth, the MODIS Team has added 70GB/day for Level 3 Ocean products generated from combined Terra and Aqua Level 2 input products.

Perhaps, the single most important point from the standpoint of MODIS Ocean science is that reducing product volume from 188GB per day to 29GB per day is not feasible in a system without a massive on-demand processing capability and the ability to subset parameters from a multi-parameter file. At 29GB per day Level 2 Ocean products cannot be archived and these Level 2 Ocean products with 1 kilometer resolution are critical for the understanding of ocean processes particularly in coastal environments. They are also vital for validation experiments, reveal fine scale phenomena not easily discernable at 5km, and are crucial inputs to higher level products that other investigators outside the MODIS team may generate. Further, at 29GB per day, MODIS won't be able to store many of the Level 3 global binned and mapped products, let alone the 1km Level 2 products unless we go to a rolling archive. In a rolling archive, all ocean products would be available for about 3 months, any user who could both predict what was needed and could get all the products before they were overwritten would be satisfied with this arrangement. Given the current distribution bandwidths this appears unlikely. Other issues include the limited availability of a Level 2 or Level 3 product for a given date (this depends on if it is in the rolling archive and if not when it would be reprocessed) and ultimately not being able to archive an Ocean time series.

MODIS Land Products

The main differences in what the MODIS Land team currently needs for the Land products in relation to the estimate that was made for the '96 baseline can be broken down into four areas:

1) estimation errors,
2) removal of old products that are no longer needed,
3) inclusion of intermediate products, and
4) the addition of new products.

First, there were a number of errors in the estimation of the product volume because the estimates were made early in the development cycle.

Second, it was determined after reviewing the Land product list, that a number of products were no longer needed. This included a few summary and monthly products.

Third, a number of new intermediate products have been included that will be very beneficial to the user community. These intermediate products have been used extensively within the Land science team to speed the development of the Level 3 algorithms. The Land Science Team believes the archive and distribution of these products will allow the land community to develop advanced science algorithms quickly and efficiently.

Fourth, two new science team members were added to the MODLAND science team after the original '96 baseline. Resources for their products, Weekly Land Surface Reflectance, Vegetation Cover Conversion and Vegetation Continuous Fields, are needed.

Finally in the '96 baseline, no estimates were given for a number of post-launch products, including Snow and Sea-ice Albedo and Burn Scar. These estimates are now included. The polar community has strongly requested that the snow products be available in polar as well as the equatorial grid to service their user community. The estimates have been updated to include these polar products.

MODIS Processing Load Growth

Processing loads have grown between the 2/96 baseline and the current request due primarily to the desire on the part of all instrument teams to be able to complete reprocessing of a year of their data within 3 to 4 months. A reprocessing rate of 4x allows an instrument team to quickly produce products and assess the quality of a time series of products. Without the higher reprocessing factor, MODIS is unlikely to have a high quality time series of products by the end of the Terra mission in 2005. Delays in developing a good time series will result in increased cost if the science team is still supported to work with the instrument products. Without this support, important science questions about global change and global climate systems will go unanswered.

What MODIS has done to reduce product volume and what could be tried next?

Over the course of developing software for the MODIS science algorithms, the MODIS science team has taken a number of steps to achieve the '96 Baseline. A number of times when the software for an algorithm came in over the processing load baseline, the scientists were asked to look for optimizations in the software that could save CPU cycles without compromising the accuracy of the products. In many cases, the redelivered software came in close to or below the baseline. To save archive storage some intermediate products were dropped from the DAAC archive/distribution list.

The MODIS land team limited the geographic extent of the land products to regions where the products were useful, i.e., non-ocean regions for most products, polar regions for sea-ice. The MODIS ocean team reduced the volume required to store their Level 3 products by converting from real values at each pixel to scaled integers. If necessary, they may limit products shipped to the Goddard DAAC based on Q/A metadata and limit Level 2 products to an archived set of granules and a 6-month rolling archive of the remaining Level 2 granules.

With the 96 Baseline the land post-launch products, which were of a more experimental nature e.g. burned area were removed from the operational production chain into the SCF environment for development. The science team continues to look for savings within the overall framework of processing, archiving and distribution. As noted earlier, ocean product volume grew when the ECS was unable to support ordering and distribution of individual parameters from a multi-parameter ocean product, such as one of the three Ocean Color products. If the ECS can demonstrate the capability to extract individual parameters and data common to all parameters from a multi-parameter file, MODIS will consider combining ocean products to reduce the volume. We are also considering using data compression within HDF-EOS to reduce the cost of distributing products to the end-user and ease the burden of storing products at the end-user's workstation. However, since hardware compression is used in the archive, it is unlikely that software compression will result in significant volume reduction though cost savings on distribution might be applied to increasing storage in the archive.